

What is Claimed is:

- Sub
Al
1. A knowledge-based system adapted to provide a recommendation tailored to a consumer, the system comprising:
 - a knowledge base containing historical data;
 - rule extraction means for extracting a ruleset from the knowledge base;
 - a rules database for holding the ruleset;
 - codifying means for codifying the requirements of the consumer;
 - recommendation means for applying the ruleset to the codified consumer requirements and generating the recommendation accordingly;
 - rule induction means for providing learning inputs to the knowledge base from a plurality of human experts as they advise and make real-life recommendations to actual or imaginary consumers, the learning inputs reflecting the recommendations made by the experts and the requirements of the consumers that they have advised; and
 - update means for running the rule extraction means on the knowledge base to refresh the rules database by extracting an updated ruleset from the knowledge base for application by the recommendation means to the requirements of future consumers.
 2. The system of Claim 1, wherein the update means operates periodically and the rule induction means provides learning inputs to the knowledge base at least as frequently as the update means operates to extract an updated ruleset from the knowledge base.
 3. The system of Claim 2, wherein the rule induction means operates continuously to provide learning inputs as they are made available by the plurality of human experts.
 4. The system of Claim 1, wherein the update means is adapted to update the knowledge base with changing details of recommendations that can be made.

5. The system of Claim 1, wherein the codifying means includes means for codifying the perceived needs of the consumer and means for codifying the circumstances of the consumer.

6. The system of Claim 1, wherein the rule induction means generates fuzzy sets, and the recommendation means applies fuzzy rules.

7. The system of Claim 6, wherein the rule extraction means implements Lozowski's algorithm.

8. The system of Claim 7, wherein Lozowski's algorithm is modified.

9. The system of Claim 8, further comprising attribute vector generating means for generating attribute vectors incrementally.

10. The system of Claim 9, wherein the attribute vector generating means is arranged such that only the last attribute vector generated is kept.

11. The system of Claim 10, wherein there is no storage of attribute vectors.

12. The system of Claim 9, further comprising means for fusing creation of attribute vectors with the evaluation of T-Norm sets.

13. The system of Claim 9, further comprising incrementing means for incrementing the attribute vectors by generating a first attribute vector that contains the first fuzzy set for each attribute, and generating the next attribute vector by selecting the next fuzzy set of the first attribute in the first attribute vector.

Al
Conf.

CONFIDENTIAL

14. The system of Claim 13, wherein the incrementing means is arranged such that if the first attribute contains no more fuzzy sets to select, the next attribute that contains more fuzzy sets to select is selected, the next fuzzy set of the selected attribute is selected, and the first fuzzy set of each lesser attribute than the selected attribute is selected.

15. The system of Claim 12, further comprising means for calculating a maximum T-Norm value while T-Norms are being generated.

16. The system of Claim 15, further comprising means for fusing generation of T-Norms and S-Norms.

17. The system of Claim 8, further comprising pruning means for pruning an attribute tree by eliminating attributes that play no part in rule-building.

18. The system of Claim 17, wherein the pruning means is arranged to mark a fuzzy set that returns zero for a current dataset example, and to omit evaluation of any attribute vector that includes the marked fuzzy set.

19. The system of Claim 18, wherein the pruning means is arranged to prune an attribute vector at class i by incrementing to the next fuzzy set value for the i -th digit of the vector while resetting any lesser digits to 0.

20. The system of Claim 19, wherein the pruning means is arranged such that if the i -th digit of the attribute vector contains no more fuzzy sets, the $i+1$ -th digit is incremented.

21. The system of Claim 1, further comprising storage means for storing a consumer's details for later recall.

22. The system of Claim 21, wherein the storage means stores the consumer's details for later recall in providing a future recommendation.

23. The system of Claim 22, wherein the storage means stores the consumer's details for later recall in completing a recommendation, and is associated with means for suspending input of consumer requirements.

24. The system of Claim 1, further comprising communications means for obtaining advice from a remote human adviser.

25. The system of Claim 24, wherein the communications means comprises a video conference link between the consumer and the adviser.

26. The system of Claim 1, wherein the system comprises a distributed system.

27. The system of Claim 26, further comprising a server holding the ruleset and data on consumers and on possible recommendations, a consumer terminal providing an online interface with the server, and a plurality of expert terminals operable by the plurality of human experts, each expert terminal including means for storing recommendations made by an expert and the requirements of consumers that that expert has advised, and means for providing that stored data to the server for use in updating the ruleset.

28. The system of Claim 27, wherein the consumer terminal is a home PC or a kiosk, booth, ATM or other terminal in a financial advice establishment.

29. The system of Claim 27, further comprising an application server enabling consumer terminals and expert terminals to interact with the server online, the application server providing an online interface to the server for the consumer terminals and the expert terminals.

30. The system of Claim 29, wherein the application server runs server-side web applications, a first web application responding to calls from a consumer website, and a second web application allowing access to the server by authorized expert terminals.

31. The system of Claim 30, wherein the server-side web applications are Java Servlets.

32. The system of Claim 31, wherein the ruleset and data on consumers and on possible recommendations are stored at the server as XML (Extensible Markup Language) documents.

33. The system of Claim 1, wherein the recommendation means includes an agent that applies the rules to the codified consumer requirements to generate the recommendation.

34. The system of Claim 33, wherein the agent includes questioning means for obtaining the consumer's requirements, the questioning means asking a sequence of questions and including means for adapting later questions in the sequence in accordance with answers given to earlier questions in the sequence.

35. A method of building a knowledge-based system for providing a recommendation tailored to a consumer, the system operating by extracting rules from a knowledge base and applying the extracted rules to codified consumer requirements to generate the recommendation accordingly; the method comprising:

providing learning inputs to the knowledge base from a plurality of human experts as they advise and make real-life recommendations to actual or imaginary consumers, the inputs reflecting the recommendations made by the experts and the requirements of the consumers that they have advised; and

after learning inputs have been provided to the knowledge base, extracting updated rules from the knowledge base for use in generating recommendations tailored to the requirements of future consumers.

36. The method of Claim 35, wherein updating operates periodically and learning inputs are provided to the knowledge base at least as frequently as the updating operates to extract updated rules from the knowledge base.

37. The method of Claim 36, wherein the learning inputs are continuously provided as they are made available by the plurality of human experts.

38. The method of Claim 35, further comprising updating the knowledge base with changing details of recommendations that can be made.

39. The method of Claim 35, wherein the codified consumer requirements include the perceived needs of the consumer and the circumstances of the consumer.

40. The method of Claim 35, further comprising employing Lozowski's algorithm to extract updated rules from the knowledge base.

41. The method of Claim 40, wherein Lozowski's algorithm is modified.

42. The method of Claim 41, further comprising generating attribute vectors incrementally.

43. The method of Claim 42, wherein only the last attribute vector generated is kept.

44. The method of Claim 43, wherein there is no storage of attribute vectors.

45. The method of Claim 44, further comprising fusing creation of attribute vectors with evaluation of T-Norm sets.

46. The method of Claim 45, wherein the attribute vectors are incremented by generating a first attribute vector that contains the first fuzzy set for each attribute, and generating the next attribute vector by selecting the next fuzzy set of the first attribute in the first attribute vector.

47. The method of Claim 46 wherein, if the first attribute contains no more fuzzy sets to select, the method comprises selecting the next attribute that contains more fuzzy sets to select, selecting the next fuzzy set of the selected attribute, and selecting the first fuzzy set of each lesser attribute than the selected attribute.

48. The method of Claim 47, further comprising calculating a maximum T-Norm value while T-Norms are being generated.

49. The method of Claim 48, wherein T-Norm and S-Norm generating steps are fused together.

50. The method of Claim 49, further comprising pruning an attribute tree by eliminating attributes that play no part in rule-building.

51. The method of Claim 50, wherein pruning is effected by marking a fuzzy set that returns zero for a current dataset example, and omitting evaluation of any attribute vector that includes the marked fuzzy set.

52. The method of Claim 51, wherein an attribute vector is pruned at class i by incrementing to the next fuzzy set value for the i -th digit of the vector while resetting any lesser digits to 0.

53. The method of Claim 52 wherein, if the i -th digit contains no more fuzzy sets, the $i+1$ -th digit is incremented.

54. The method of Claim 53, wherein the knowledge base is established by liaising with experts to:

determine the primary parameters of a consumer profile;
 codify the possible values of those parameters;
 generate and run a plurality of case studies; and
 generate an initial set of rules by applying the parameters to the case studies.

55. The method of Claim 54, wherein each parameter to be taken into account is broken down into fuzzy sets and weighted.

56. The method of Claim 55, wherein a set of parameters are contained in an XML (Extensible Markup Language) DTD (document type definition).

57. The method of Claim 56, further comprising mapping the parameters to recommendations resulting from the case studies to generate the initial set of rules.

58. The method of Claim 57, wherein the initial set of rules is generated by running Lozowski's algorithm.

59. The method of Claim 58, further comprising holding on a server the rules and data on consumers and on possible recommendations, storing on a remote terminal recommendations made by an expert and the requirements of consumers that that expert has advised, and providing that stored data to the server for use in updating the rules.

60. A method of operating a knowledge-based system for providing a recommendation tailored to a consumer, the method comprising:
 extracting rules from a knowledge base and applying the extracted rules to codified consumer requirements to generate the recommendation accordingly;
 providing learning inputs to the knowledge base from a plurality of human experts as they advise and make real-life recommendations to actual or imaginary consumers,

the inputs reflecting the recommendations made by the experts and the requirements of the consumers that they have advised; and

after learning inputs have been provided to the knowledge base, extracting updated rules from the knowledge base for use in generating recommendations tailored to the requirements of future consumers.

61. The method of Claim 60, further comprising storing a consumer's details for later recall.

62. The method of Claim 61, wherein the consumer's details are stored and later recalled when providing a future recommendation.

63. The method of Claim 62, wherein the consumer's details are stored and later recalled when completing a recommendation, input of consumer requirements being suspended in the meantime.

64. The method of Claim 63, further comprising the consumer obtaining advice from a remote human adviser.

65. The method of Claim 64, further comprising obtaining the advice over a video conference link.

66. The method of Claim 65, wherein the consumer enters consumer requirements via a consumer terminal being a home PC or a kiosk, booth, ATM or other terminal in a financial advice establishment.

67. The method of Claim 66, wherein the consumer enters consumer requirements in response to a sequence of questions, and wherein later questions in the sequence are adapted in accordance with answers given to earlier questions in the sequence.